

Internet Engineering Task Force
INTERNET-DRAFT
draft-ietf-idmr-msnip-01.txt

IDMR WG
Bill Fenner/AT&T
Hugh Holbrook/Cisco
Isidor Kouvelas/Cisco
21 November 2001
Expires: May 2002

Multicast Source Notification of Interest Protocol (MSNIP)

Status of this Document

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

This document is a product of the IETF IDMR WG. Comments should be addressed to the authors, or the WG's mailing list at
pim@catarina.usc.edu.

Abstract

This document discusses the Multicast Source Interest Notification Protocol (MSNIP). MSNIP is an extension to IGMPv3 [1] that provides membership notification services for sources of multicast traffic.

INTERNET-DRAFT

Expires: May 2002

November 2001

Table of Contents

1. Introduction.	3
2. Routing Protocol Support.	3
3. API for Requesting Membership Notification.	4
4. MSNIP Managed Address Range Negotiation	5
4.1. Router Coordination.	6
4.1.1. MSNIP Capability Discovery Option	6
4.1.2. SSM Range Discovery Option.	7
4.2. Communicating Range to Source Systems.	7
4.2.1. MSNIP Range Map Messages.	8
5. Requesting and Receiving Notifications.	9
5.1. Host Interest Solicitation	10
5.2. Router Receiver Membership Reports	10
6. Routerless Operation.	11
7. Application Notification.	12
8. Router Processing	14
9. Message Formats	16
9.1. Range Map Packet	16
9.2. Interest Solicitation Packet	18
9.3. Receiver Membership Report Packet.	19
10. Constants Timers and Default Values.	20
11. Todo list...	21
12. Acknowledgements	21
13. Authors' Addresses	21
14. References	22

1. Introduction

The Multicast Source Notification of Interest Protocol (MSNIP) is an extension to version 3 of the Internet Group Membership Protocol (IGMPv3 [1]). MSNIP operates between multicast sources and their first-hop routers to provide information on the presence of receivers to the source systems. Using the services offered by MSNIP an application on an IP system wishing to source multicast data can register to be notified when receivers join and leave the session. This enables multicast sources to avoid the work of transmitting packets onto their first-hop link when there are no joined receivers.

A common scenario where MSNIP may be useful is one where there is a multicast server offering a large pool of potential flows that map onto separate multicast destination addresses but receivers exist only for a small subset of the flows. If the source were to continuously transmit data for all the flows that could potentially have receivers, significant resources would be wasted in the system itself as well as the first-hop link. Using a higher level control protocol to determine the existence of receivers for particular flows may not be practical as there may be many potential receivers in each active session.

Information on which multicast destination addresses have receivers for a particular sender is typically available to the multicast routing protocol on the first hop router for a source. MSNIP uses this information to notify the application in the sending system of when it should start or stop transmitting. This is achieved without any destination address specific state on the first-hop router for potential sources without receivers.

2. Routing Protocol Support

For reasons described in this section, MSNIP only supports transmission control for applications that use multicast destination addresses that are routed using Source Specific Multicast (SSM).

Many currently deployed multicast routing protocols, require data from an active source to be propagated past the first-hop router before

information on the existence of receivers becomes available on the first-hop. In addition, such protocols require that this activity is repeated periodically to maintain source liveness state on remote routers. All dense-mode protocols fall under this category as well as sparse-mode protocols that use shared trees for source discovery (such as PIM-SM [3]). In order to provide receiver interest notification for such protocols, the default mode of operation would require that the source IP system periodically transmits on all potential destination addresses and the first-hop routers prune the traffic back. Such a

Fenner/Holbrook/Kouvelas

Section 2. [Page 3]

INTERNET-DRAFT

Expires: May 2002

November 2001

flood-and-prune behaviour on the first-hop link significantly diminishes the benefits of managing source transmission.

In contrast, with source-specific sparse-mode protocols such as PIM-SSM [3] availability of receiver membership information on the first-hop routers is independent of data transmission. Such protocols use an external mechanism for source discovery (like source-specific IGMPv3 membership reports) to build source-specific multicast trees.

Clearly these two classes of routing protocols require different handling for the problem MSNIP is trying to solve. In addition the second type covers the majority of applications that MSNIP is targeted at. MSNIP avoids the extra complication in supporting routing protocols that require a flood and prune behaviour.

3. API for Requesting Membership Notification

Applications within an IP system that wish to source multicast packets and are interested in being notified on the existence of receivers must register with the IP layer of the system. MSNIP requires that within the IP system, there is (at least conceptually) an Application Programming Interface or API that can be used to register with the IP layer for such notifications. A system's IP API must support the following operation or any logical equivalent:

```
IPMulticastsSourceRegister (socket, source-address, multicast-address)
IPMulticastsSourceDeregister (socket, source-address, multicast-address)
```

In addition the application must provide the following API for receiving notifications from the IP system:

```
IPMulticastSourceStart (socket, source-address, multicast-address)
IPMulticastSourceStop (socket, source-address, multicast-address)
```

where:

socket

is an implementation-specific parameter used to distinguish among different requesting entities (e.g., programs or processes) within the system; the socket parameter of BSD Unix system calls is a specific example.

source-address

is the IP unicast source address that the application wishes to use in transmitting data to the specified multicast address. The specified address must be one of the IP addresses associated with the network interfaces of the IP system. Note that an interface in an IP system may be associated with more than one IP addresses. An

Fenner/Holbrook/Kouvelas

Section 3. [Page 4]

INTERNET-DRAFT

Expires: May 2002

November 2001

implementation may allow a special "unspecified" value to be passed as the source-address parameter, in which case the request would apply to the "primary" IP address of the "primary" or "default" interface of the system (perhaps established by system configuration). If transmission to the same multicast address is desired using more than one source IP address, IPMulticastSourceRegister is invoked separately for each desired source address.

multicast-address

is the IP multicast destination address to which the request pertains. If the application wishes to transmit data to more than one multicast address for a given source address, IPMulticastSourceRegister is invoked separately for each desired multicast address.

Applications wishing to use MSNIP to control their multicast data transmission to destination G from source address S operate as follows.

Initially the application contacts the IP system to obtain the socket handle for use on all subsequent interactions. The application invokes IPMulticastSourceRegister for the desired S and G and then waits without transmitting any packets for the IP system to notify that receivers for the session exist.

If and when the IP system notifies the application that receivers exist using the `IPMulticastSourceStart` call, the application may start transmitting data. After the application has been notified to send, if all receivers for the session leave, the IP system will notify the application using the `IPMulticastSourceStop` call. At this point the application should stop transmitting data until it is notified again that receivers have joined through another `IPMulticastSourceStart` call.

When the application no longer wishes to transmit data it should invoke the `IPMulticastsSourceDeregister` call to let the IP system know that it is no longer interested in notifications for this source and destination. The `IPMulticastsSourceDeregister` call should be implicit in the teardown of the associated socket state.

4. MSNIP Managed Address Range Negotiation

With current multicast deployment in the Internet, different multicast routing protocols coexist and operate under separate parts of the multicast address space. Multicast routers are consistently configured with information that maps specific multicast address ranges to multicast routing protocols. Part of this configuration describes the

Fenner/Holbrook/Kouvelas

Section 4. [Page 5]

INTERNET-DRAFT

Expires: May 2002

November 2001

subset of the address space that is used by source-specific multicast (SSM) [4]. As described in section 2 MSNIP only tries to control application transmission within the SSM address range.

It is desirable for applications within an IP system that supports MSNIP to have a consistent API for multicast transmission that does require the application to be aware of the SSM address range. MSNIP supports this by allowing applications to use the API described in section 3 for multicast destination addresses that are outside its operating range. When an application registers for notifications for a destination address that is not managed by MSNIP it is immediately notified to start transmitting. This complies with the default behaviour of IP multicast without MSNIP support which forces multicast applications to assume that there are multicast receivers present in the network.

4.1. Router Coordination

In order for MSNIP to operate on a shared link where more than one


```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
Destination-Prefix-1                                     | Mask-Len-2 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Destination-Prefix-2                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                                                               ...                               |

```

Length

The length of the SSM Range Discovery option is variable and is equal to five times the number of destination ranges present in the option.

Destination-Prefix-n

The multicast destination address prefix of nth range present in this option.

Mask-Len-n

The mask length for the nth address range.

A router receiving a Multicast Router Discovery message with an SSM Range Discovery Option must compare the contents of the option with the multicast address ranges in the local SSM configuration and signal any differences to the administrator.

4.2. Communicating Range to Source Systems

When an application in an IP system uses the MSNIP API to register for notification, the IP system must behave differently depending on

whether or not the destination address for which the application registered is operating under SSM. For SSM channels, the IP system should only instruct the application to transmit when there are receivers for the multicast destination. For non-SSM destination addresses the IP system will not be able to determine if there are receivers and should immediately instruct the application to transmit.

To be able to differentiate between SSM and non-SSM multicast

destination addresses a source IP system must know the SSM range which is configured on the first-hop routers. One method for an IP system to discover the SSM range would be to listen to Multicast Router Advertisement messages and use the information in the SSM Range option. However, as currently specified, IGMP Multicast Router Discovery messages are sent to the ALL-Routers (224.0.0.2) multicast destination address to which host IP systems do not normally subscribe.

Three options are under consideration for the mechanism to distribute the MSNIP managed range (SSM range) to host IP systems:

- o Make hosts listen to the ALL-Routers multicast destination for IGMP Multicast Router Discovery messages.
- o Use a separate MSNIP control message and nominate a router as responsible for communicating the range to hosts. This option is presented in more detail below using the IGMP querier as the router responsible for distributing the MSNIP / SSM range.
- o Do not use the IGMP Multicast Router Discovery SSM Range option to distribute the information between routers. Instead use an MSNIP message that hosts receive as well. Note that such a separate mechanism would not have the size limitations of the router discovery option.

4.2.1. MSNIP Range Map Messages

This section describes the option of using a separate MSNIP control message for communicating the SSM multicast address range to host IP systems. Communicating the range is left up to the IGMP querying router. The querier periodically multicasts a MSNIP Range Map message containing the definition of the address ranges over which MSNIP operates. The destination of the Range Map message is the ALL-SYSTEMS multicast address. The Range Map message is sent every [Range Map Interval] seconds. The message also contains a holdtime which is set to [Range Map Holdtime] and instructs IP systems to maintain the range mapping state for the specified holdtime.

In addition to the periodic transmission, triggered Range Map messages

are sent when either:

- o the IGMP querier on a link receives an Interest Solicitation message (described in section 5.1) from an IP system that was not previously registered with MSNIP or was registered with a different GenID (see section 9.2).
- o the configured SSM address range on the querier changes.

When either of these two events occur the querier initiates transmission of a set of [Robustness Variable] Range Map messages.

Upon receipt of a Range Map message, an IP system builds or updates a set of range records as follows. For each multicast address range present in the message, the system either creates or updates a record of the form:

(interface, range prefix, range mask)

Where interface is the interface the Range Map message was received on and prefix and mask are the definition of the address range. If range records which were created by a previous Range Map message received on this interface are not present in the current message, these records are deleted.

In addition to the address range records, the IP system maintains a holdtime timer associated with the interface which is initialised to the holdtime in the received message. If the timer expires before the receipt of the next Range Map message, all multicast address range records related to the interface are deleted.

5. Requesting and Receiving Notifications

Like IGMP, MSNIP is an asymmetric protocol specifying different behaviours for systems wishing to source traffic and for multicast routers. Host IP systems multicast Host Interest Solicitation messages to register for notification with their first-hop routers. Routers unicast Router Receiver Membership Reports to IP system to notify them of the arrival of the first or departure of the last receivers for a session. Note that a system may perform at the same time both of the above functions. An example is a router that wishes to source traffic.

5.1. Host Interest Solicitation

Source systems that wish to be managed by MSNIP periodically transmit an Interest Solicitation message. This message is multicast with a multicast destination address of ALL_IGMPv3_ROUTERS (224.0.0.22) and is transmitted every [Interest Solicitation Interval] seconds. The Interest Solicitation message contains a holdtime which is set to [Interest Solicitation Holdtime] and instructs the multicast first-hop routers to maintain MSNIP state for this IP system for the specified period. A generation ID is also included in the Interest Solicitation message to provide support for routers to detect IP system crashes (see section 9.2).

When an IP system first comes up it transmits [Robustness Variable] Interest Solicitation messages randomly spaced over [Initial Interest Solicitation Interval] seconds.

All MSNIP capable routers that receive an Interest Solicitation message from an IP system, maintain a system interest record of the form:

(IP system address, holdtime timer)

Each time an Interest Solicitation message is received from the IP system, the holdtime timer is reset to the holdtime in the received message. In addition the router responds to the solicitation message with a Receiver Membership Report message described in section 5.2. The message contains a TRANSMIT record for each of the multicast destination addresses within the MSNIP managed range for which the routing protocol indicates there are receivers for this source system.

When the holdtime timer of a specific system interest record expires, the record is deleted.

5.2. Router Receiver Membership Reports

Receiver Membership Report messages are used by routers to communicate the receiver membership status of particular multicast destination addresses to a specific IP system. Each message contains a list of transmission control records of either TRANSMIT or HOLD type that instruct a system to respectively start or stop sending traffic on this link to the specified multicast destination address. Receiver Membership Report messages are unicast to the target system.

In addition to the receipt of an Interest Solicitation message, routers send unsolicited Receiver Membership Reports to IP systems when the receiver membership status reported by the multicast routing

Fenner/Holbrook/Kouvelas

Section 5.2. [Page 10]

INTERNET-DRAFT

Expires: May 2002

November 2001

protocol changes for a specific source and multicast destination. Such reports are only sent if the destination address is managed by MSNIP and the router has a system interest record created by a previously received Interest Solicitation message with a IP system address equal to the source address. If the source destination pair satisfy these conditions then [Robustness Variable] Receiver Membership Reports are sent out within [Unsolicited Membership Report Interval] seconds. If during the [Unsolicited Membership Report Interval] an additional membership change occurs for the same destination address and source system, transmission of any related pending Receiver Membership Report messages is cancelled and a new set of [Robustness Variable] messages is scheduled.

When an IP system receives a Receiver Membership Report message, for each of the TRANSMIT records listed in the message it creates or updates a transmission record of the form:

(router address, source address, multicast address, holdtime timer)

The router address is obtained from the source address on the IP header of the received message. The source address is obtained from the destination address in the of the IP header. The holdtime timer is set to [Interest Solicitation Holdtime] seconds.

For each HOLD record present in the message, the system deletes the matching previously created transmission record from its state.

Note that creation and deletion of transmission records in an IP systems state may cause local applications to be notified to start and stop transmitting data (see section 7).

6. Routerless Operation

As defined in this specification MSNIP provides receiver membership notification services for multicast networks operating an SSM routing protocol. The protocol operates between the source IP system and the first-hop routers. Although not obvious, local receivers on the same link as the multicast source do not require special handling. Local receivers use the usual process of joining the SSM channel through IGMPv3 source specific joins. IGMP makes the local membership

information available to the routing protocol on all routers on the link. The routing protocol is then responsible for electing the router that will be responsible for acting on behalf of the local receivers (in PIM this is the Designated Router or the PIM Assert winner). Normal operation then results in MSNIP being notified and the source signalled to start transmitting.

A special case that is not handled by default is that of a link not connected to a routed multicast network. On a link with only senders and

Fenner/Holbrook/Kouvelas

Section 6. [Page 11]

INTERNET-DRAFT

Expires: May 2002

November 2001

receivers but no routers MSNIP capable sources do not have a mechanism for being notified about the existence of local receivers. However, since there is no router to send out Range Map messages, IP systems assume that there are no MSNIP managed address ranges and all applications default to transmitting immediately. Therefore expected behaviour (without MSNIP) is preserved.

In order to prevent source flooding in a routerless link when there are no local receivers for the data, MSNIP requires that one of the IP systems on the link acts as an MSNIP server. This server must implement the router side of the IGMPv3 and MSNIP protocols. The MSNIP server must be configured with a multicast address range that is to be managed which will then be advertised in the Range Map messages. When IGMPv3 Source Specific reports are received for sources on the link, the IGMP component in the server must notify the MSNIP component. If the multicast destination address for which the report was received is a managed address then MSNIP can perform its usual functions to control the source.

7. Application Notification

This section describes the relation between protocol events and notifications to source applications within an IP system. The state machine below is specific to each source address of the IP system and each multicast destination address. The initial state is the No Info state.

```
+-----+
| Figures omitted from text version |
+-----+
```

Figure 1: Per source-address (S) and multicast destination address (G) state machine at an IP system

Fenner/Holbrook/Kouvelas

Section 7. [Page 12]

INTERNET-DRAFT

Expires: May 2002

November 2001

In tabular form, the state-machine is:

+-----+-----+-----+-----+-----+-----+						
		Event				
+-----+-----+-----+-----+-----+-----+						
Prev State	New	Start	Stop	Recv	Recv last	
	Register	Manage	Manage	TRANSMIT	HOLD or	
					timeout	
+-----+-----+-----+-----+-----+-----+						
	Start new	-> Hold	-	-	-	
No Info		Stop ALL				
		registered				
+-----+-----+-----+-----+-----+-----+						
	-	-	-> No Info	->	-	
Hold				Transmit		
			Stop ALL	Start ALL		
			registered	registered		
+-----+-----+-----+-----+-----+-----+						
	Start new	-	-> No Info	-	-> Hold	
Transmit					Stop ALL	
					registered	
+-----+-----+-----+-----+-----+-----+						

The events in state machine above have the following meaning:

New register

A new application has registered through a call to IPMulticastsSourceRegister for this S and G.

Start manage

We received a Range Map packet on the interface that S belongs to that changed the status of G from a non-managed to a MSNIP managed destination address.

Stop manage

We received a Range Map packet on the interface that S belongs to that changed the status of G from a MSNIP managed to a non-managed destination address or the mapping state that caused this destination address to be managed expired.

Receive TRANSMIT

We received a Receiver Membership Report with S as the IP destination address that contains a TRANSMIT record for G.

Fenner/Holbrook/Kouvelas

Section 7. [Page 13]

INTERNET-DRAFT

Expires: May 2002

November 2001

Receive last HOLD or timeout

We either received a Receiver Membership Report with S as the IP destination address that contains a HOLD record for G or a TRANSMIT record for S and G expired and there are no other TRANSMIT records for S and G. This means that the last router that was reporting receivers no longer does so and there are no routers left wishing to receive traffic from this S to destination address G.

The state machine actions have the following meaning:

Start new

Send an IPMulticastSourceStart notification to the application that just registered for this S and G.

Start ALL registered

Send an IPMulticastSourceStart notification to all applications that are registered for this S and G.

Stop ALL registered
Send an IPMulticastSourceStop notification to all applications that
are registered for this S and G.

8. Router Processing

This section describes the per-source system tracking state machine
operated by each first-hop router. The initial state is No Info.

+-----+
| Figures omitted from text version |
+-----+

Figure 2: Per IP source system (S) state machine at a router

In tabular form, the state-machine is:

+-----+				
	Event			
	+-----+			
Prev State	Receive	HIS	Receivers	Receivers
	HIS	timeout	for new	of G leave
			destination	
			G	
+-----+				
	->	-	-	-
	Tracking			
	Set HT to			
Not	message			

tracking	holdtime				
	Send ALL				
	existing				
	TRANSMITs				
+-----+-----+-----+-----+-----+					
Tracking	Set HT to	-> Not	Send	Send HOLD	
	message	tracking	TRANSMIT	for G	
	holdtime		for G		
+-----+-----+-----+-----+-----+					

The events in state machine above have the following meaning:

Receive HIS

The router has received a Host Interest Solicitation from S.

HIS timeout

The holdtime timer (HT) associated with S has expired.

Receivers for new destination G

The routing protocol has informed MSNIP that it now has receivers for the MSNIP managed destination address G and source IP system S.

Receivers of G leave

The routing protocol has informed MSNIP that all receivers for the MSNIP managed destination address G and source IP system S have left the channel.

The state machine actions have the following meaning:

Fenner/Holbrook/Kouvelas

Section 8. [Page 15]

INTERNET-DRAFT

Expires: May 2002

November 2001

Set HT to message holdtime

The holdtime timer associated with S is restarted to the value of the holdtime in the received Host Interest Solicitation message.

Send ALL existing TRANSMITs

The router builds and transmits Receiver Membership Reports to S

that contain a TRANSMIT record for each of the MSNIP managed destination addresses that have receivers for S.

Send TRANSMIT for G

The router builds and transmits a Receiver Membership Report to S that contains a TRANSMIT record for the destination address G.

Send HOLD for G

The router builds and transmits a Receiver Membership Report to S that contains a HOLD record for the destination address G.

9. Message Formats

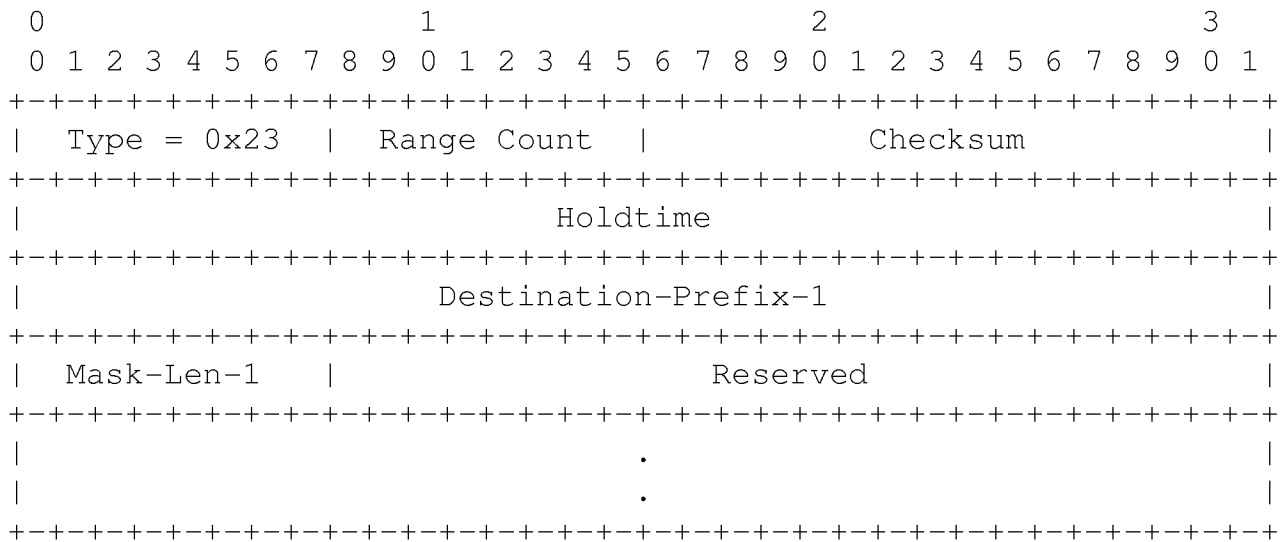
Like all other IGMP messages, MSNIP messages are encapsulated in IPv4 datagrams, with an IP protocol number of 2. Every MSNIP message described in this document is sent with an IP Time-to-Live of 1, and carries an IP Router Alert option [RFC-2113] in its IP header.

There are three IGMP message types of concern to the MSNIP protocol described in this document:

+-----+-----+		+-----+
Type Number (hex)	Message Name	
+-----+-----+		+-----+
0x23	Range Map	
+-----+-----+		+-----+
0x24	Interest Solicitation	
+-----+-----+		+-----+
0x25	Receiver Membership Report	
+-----+-----+		+-----+

9.1. Range Map Packet

A Range Map packet is periodically multicast by the IGMP querier to declare the multicast destination address ranges managed by MSNIP. The Range Map message is multicast with a destination address of ALL_SYSTEMS



Range Count

The number of multicast destination address range records present in this message. Note that the actual maximum number of ranges that can be reported is limited by the maximum size of an IP packet. As each Range Map packet replaces the mapping at a receiving system, a router cannot split the range mapping into more than one Range Map packets.

Checksum

The Checksum is the 16-bit one's complement of the one's complement sum of the whole MSNIP message (the entire IP payload). For computing the checksum, the Checksum field is set to zero. When receiving packets, the checksum MUST be verified before processing a packet.

Holdtime

The amount of time a receiving system must keep the range map state alive, in seconds.

Destination-Prefix-1

The destination address range prefix of the first range record present in this message.

INTERNET-DRAFT

Expires: May 2002

November 2001

Mask-Len-1

The mask length for the destination address range in the first record present in this message.

Reserved

Transmitted as zero. Ignored upon receipt.

9.2. Interest Solicitation Packet

A Interest Solicitation packet is periodically multicast by MSNIP capable systems to declare interest in Receiver Membership Reports from multicast routers on the link. The Interest Solicitation message is multicast with a destination address of ALL_IGMPv3_ROUTERS (224.0.0.22).

[illegible]

Reserved

Transmitted as zero. Ignored upon receipt.

Checksum

Calculated as for Range Map packet.

Holdtime

The amount of time a receiving router must keep the system interest state alive, in seconds.

GenID

Generation ID of the IP system. A number that is selected randomly for each of the [Robustness Variable] initial Interest Solicitation messages when the system comes up and afterwards remains fixed to the value used in the last of the initial messages throughout the

Section 9.2. [Page 18]

November 2001

A Receiver Membership Report packet is unicast by first-hop multicast routers and targeted at potential sources to inform them of the existence or not of receivers for the listed multicast destination addresses.

The number of multicast destination address records present in this message.

Calculated as for Range Map packet.

The type of the first transmission control record in this message.
Valid values are:

Record Type	Description	Value
-------------	-------------	-------

+-----+-----+-----+-----+			
TRANSMIT	Request to start transmitting to destination	1	
+-----+-----+-----+-----+			
HOLD	Request to stop transmitting to destination	2	
+-----+-----+-----+-----+			

Reserved

Transmitted as zero. Ignored upon receipt.

Fenner/Holbrook/Kouvelas

Section 9.3. [Page 19]

INTERNET-DRAFT

Expires: May 2002

November 2001

Destination-Address-1

The multicast destination address of the first record in the message.

10. Constants Timers and Default Values

Robustness Variable

The Robustness Variable allows tuning for the expected packet loss on a network. If a network is expected to be lossy, the Robustness Variable may be increased. MSNIP is robust to (Robustness Variable - 1) packet losses. The Robustness Variable MUST NOT be zero, and SHOULD NOT be one. Default: 2

Range Map Interval

The interval used by the IGMP querier between transmissions of Range Map messages. Default: 60 secs

Range Map Holdtime

The interval inserted in Range Map messages that indicates to systems how long they should use the included mapping information before they time it out. This MUST be ((the Robustness Variable) times (the Range Map Interval) plus (one second)).

Interest Solicitation Interval

The interval used by MSNIP capable systems between transmissions of Interest Solicitation messages. Default: 60 secs

Interest Solicitation Holdtime

The interval inserted in Interest Solicitation messages by systems to instruct routers how long they should maintain system interest state for. This MUST be ((the Robustness Variable) times (the Interest Solicitation Interval) plus (one second)).

Initial Interest Solicitation Interval

The interval used by systems to send out the initial Interest Solicitation messages when they first come up. Default: 1 second.

Unsolicited Membership Report Interval

The interval used by routers to send out a set of Membership Report messages when the receiver membership changes for a specific system. Default: 1 second.

Fenner/Holbrook/Kouvelas

Section 10. [Page 20]

INTERNET-DRAFT

Expires: May 2002

November 2001

11. Todo list...

- o Add security considerations section.
- o Cover IPv6.
- o If application ignores or does not ask for notification in managed range host OS should filter packets.
- o Maybe provide masks for registration / notification API.
- o Case where host and app starts before MSNIP range is available.
- o When we hear out-of-order IGMP query resend interest registration.
- o Only send interest registration when application is interested. This is not possible if we do host filtering...

- o Maybe add API to ask the kernel for the state of a particular destination. `bool IpMulticastSourceHasInterest (socket, source-address, multicast-address).`
- o Add GenID changes to router FSM.

12. Acknowledgements

The authors would like to thank Dave Thaler and Jon Crowcroft for their contribution to this specification.

Fenner/Holbrook/Kouvelas

Section 13. [Page 21]

INTERNET-DRAFT

Expires: May 2002

November 2001

13. Authors' Addresses

Bill Fenner
AT&T Labs - Research
75 Willow Road
Menlo Park, CA 94025
fenner@research.att.com

Hugh Holbrook
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
holbrook@cisco.com

Isidor Kouvelas
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134

kouvelas@cisco.com

14. References

- [1] B. Cain, S. Deering, W. Fenner, I. Kouvelas, A. Thyagarajan, "Internet Group Management Protocol, Version 3", Work In Progress, <draft-ietf-idmr-igmp-v3-05.txt>, 2000.
- [2] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol.", RFC 2401.
- [3] B. Fenner, M. Handley, H. Holbrook, I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", Work In Progress, <draft-ietf-pim-sm-v2-new-01.txt>, 2000.
- [4] Z. Albanna, K. Almeroth, D. Meyer, "IANA Guidelines for IPv4 Multicast Address Allocation", Best Current Practices, <draft-ietf-iana-IPv4-mcast-guidelines-00.txt>, 2001.
- [5] S. Biswas, B. Haberman, "IGMP Multicast Router Discovery", Work In Progress, <draft-ietf-idmr-igmp-mrdisc-07.txt>, 2001.

Fenner/Holbrook/Kouvelas

Section 14. [Page 22]

INTERNET-DRAFT

Expires: May 2002

November 2001

